

REMARKS

This paper is submitted in reply to the Office Action dated March 8, 2006, within the three-month period for response. Reconsideration and allowance of all pending claims are respectfully requested.

In the subject Office Action, claims 29-30 were rejected under 35 U.S.C. § 101 as being directed to non-statutory subject matter. Moreover, claims 1-5, 7-8, 12-18, 20-21 and 25-30 were rejected under 35 U.S.C. § 102(b) as being anticipated by U.S. Patent No. 6,266,658 to Adya et al. Furthermore, claims 6, 9-11, 19 and 22-24 were rejected under 35 U.S.C. § 103(a) as being unpatentable over U.S. Patent No. 6,073,129 to Levine et al.

Applicants respectfully traverse the Examiner's rejections to the extent that they are maintained.

With respect to the §101 rejection, Applicants have amended claim 29 to recite a "tangible" computer readable signal bearing medium, and have canceled claim 30 without prejudice. Applicants respectfully submit that no new matter is being added by the above amendments, as the amendments are fully supported in the specification, drawings and claims as originally filed. Furthermore, Applicants submit that the aforementioned amendments adequately address the Examiner's concerns. Withdrawal of the §101 rejection is therefore respectfully requested.

Next, with respect to the art-based rejections, and more specifically to the rejection of independent claim 1, this claim generally recites a method of performing a database query. The method includes generating an access plan for the database query, estimating a percentage of the resource that is currently resident in working memory, estimating a cost for the access plan using the estimated percentage, and selectively executing the access plan based upon the estimated cost. The access plan uses at least one resource capable of being retrieved into working memory, where the resource is selected from the group consisting of a database file, a database table, an index, a temporary result set, a temporary file, and a hash table.

Applicants' invention addresses a problem experienced by conventional query optimizers where the estimated cost of an access plan as calculated by a query optimizer may not accurately represent the actual cost for executing an access plan based upon a

given runtime environment, in particular due to inaccurate estimations of input/output costs that are based upon the amount of time required to retrieve resources necessary for executing a query. Whenever a cost estimate is incorrect, a risk exists that a query optimizer will not select the optimal access plan, and thus a query will not execute with optimal efficiency.

As an example, a resource such as an index often is able to substantially accelerate the execution of a query. However, if an index is stored in persistent storage, there is an input/output cost associated with retrieving the index into working memory (e.g., RAM) before that index can be used during execution of the query. In some instances, the cost to load the index into working memory may be too excessive, and may result in the selection of another access plan that does not use the index. On the other hand, if the index has already been loaded into working memory before the query is executed (e.g., if a prior query has already loaded the index), the input/output cost associated with loading the working memory is practically eliminated, and may make an access plan that uses the index more efficient than competing access plans. Conventional query optimizers, however, are unable to ascertain whether an index has or has not already been loaded into a working memory when selecting an access plan for a query that might use such an index, and as such, are often required to assume that an index will be required to be loaded. As such, if a particular access plan that uses the index would be more efficient than competing access plans if the cost of loading the index was not taken into account, that access plan will typically not be selected even if the index was already loaded into working memory when the query is executed.

Embodiments of Applicants' invention, on the other hand, are capable of using a "retrieval status" of a resource such as a database table, file, index, or hash table to generate a more accurate input/output cost estimate for an access plan, and thus enable a query optimizer to select an optimal access plan given the current runtime conditions under which the access plan will execute. In claim 1, in particular, this retrieval status may take the form of a percentage value that represents the percentage of a required resource that is already loaded into working memory. Thus, for example, if it is determined that 45% of an index used by an access plan is already resident in working

memory, the input/output cost for loading the index may be reduced by 45%, resulting in a more accurate cost estimation for the access plan.

In rejecting claim 1, the Examiner relies on Adya, and in particular col. 1, lines 22-24, col. 3, line 40, col. 7, lines 1-7, col. 8, lines 25-30 and col. 10, lines 11-14. However, in none of these passages, nor anywhere else in Adya, can Applicants find any disclosure of a number of features recited in claim 1, e.g., "estimating a percentage of [a] resource that is current resident in working memory" and "estimating a cost for [an] access plan using the estimated percentage."

Instead, Adya is directed to a methodology for determining a set of indexes to be proposed to an administrator for the purpose of optimizing a runtime environment for a given workload. The methodology attempts to determine what indexes should be generated or maintained in storage based upon limited available storage. The goal of Adya is to recommend a set of indexes that will have the greatest positive impact on performance within the limited amount of storage space that can hold those indexes.

The specific passages relied upon by the Examiner disclose, at the most, using cost estimates for queries to rank potential indexes for recommendation to a systems administrator for a DBMS. The passages do not, in fact, even deal with executing queries; instead they deal with recommending indexes that a system administrator could manually generate or update to optimize system workload in the future. Indeed, the cost calculation disclosed at col. 8, lines 25-30 is not a cost of an access plan, it is the estimated cost of an index to determine whether that index, if created, would substantially improve the performance of the system.

Adya, however, does not ever attempt to determine whether or not an index even exists in any form of storage, much less doing so for the purpose of calculating the cost of an access plan that might use such an index. More importantly, Adya does not disclose determining a retrieval status of an index, and using that retrieval status to adjust the cost estimate for an access plan.

Given that Adya does not disclose determining a retrieval status of an index, Applicants submit that Adya cannot be interpreted as disclosing the estimation of a percentage of an index (or any other required resource) that is resident in a working

memory, or the estimation of a cost for an access plan using any estimated percentage. Accordingly, claim 1 is novel over Adya, and the rejection thereof should be withdrawn.

Applicants also submit that claim 1 is non-obvious over Adya, as there is no suggestion in the reference, or elsewhere in the prior art, of the desirability of estimating a cost for an access plan based in part of the percentage of a required resource that is resident in a working memory. Indeed, Adya provides little if any particulars on how cost estimates are even calculated for individual access plans. No evidence has been presented establishing that one of ordinary skill in the art would be motivated to modify Adya to incorporate any access plan cost estimation that is based at least in part on an estimated percentage of a resource that is resident in working memory. Accordingly, claim 1 is also non-obvious over Adya. Reconsideration and allowance of claim 1, and of claim 2 which depends therefrom, are therefore respectfully requested.

Next, with respect to independent claim 3, this claim generally recites method of optimizing a database query, which includes determining a retrieval status for a resource used by the database query, and generating an access plan for the database query using the determined retrieval status for the resource.

Claim 3 is also rejected based upon Adya, and in particular col. 3, line 27 and col. 7, lines 1-5 and 12-14. As discussed above in connection with claim 1, however, Adya does not disclose determining a retrieval status for a resource, much less doing so for the purpose of generating an access plan. As such, Adya fails to disclose "determining a retrieval status for a resource used by [a] database query," or "generating an access plan for the database query using the determined retrieval status for the resource," as required by claim 3.

Adya discloses, at the most, determining a set of indexes suitable for optimizing the performance of a DBMS in handling a set of queries, and doing so such that the set of indexes will fit within a constrained storage space. The "retrieval status" of an index is never determined in Adya, and indeed, in many instances the indexes that are being analyzed do not even exist. It is only after an administrator accepts the recommended set of indexes and creates those indexes (a point in time after the calculations described in Adya are performed) that those indexes even come into existence.

It may be that the Examiner considers the concept of determining the retrieval status of a resource to be met by a determination of whether a resource exists or not. By reciting "determining a retrieval status for a resource" however, claim 3 presumes that a resource does in fact exist. And even were this not the case, Adya does not even attempt to determine whether an index exists or not. In fact, Adya presumes that the indexes being recommended will need to be generated by an administrator after the recommendation is made (or will be generated automatically - see col. 7, lines 30-37), so whether or not these indexes already exist as of the calculations is not even germane to the calculations performed in Adya. Indexes that exist are treated no differently from indexes that do not yet exist, but could be created, so Adya does not suggest, and would have no reason to suggest, the need to determine whether an index exists.

Given that Adya fails to disclose "determining a retrieval status for a resource used by [a] database query," or "generating an access plan for the database query using the determined retrieval status for the resource," claim 3 is novel over Adya. Moreover, claim 3 is non-obvious over the reference as the Examiner has provided no objective evidence that one of ordinary skill in the art would be motivated to modify Adya to determine the retrieval status of a resource used by a database query when generating an access plan for that query. Accordingly, claim 3 is also non-obvious over Adya. Reconsideration and allowance of claim 3, as well as of claims 4-15 which depend therefrom, are therefore respectfully requested.

Next turning to independent claims 16 and 29, each of these claims recite in part program code configured to a database query by "determining a retrieval status for a resource used by the database query, and generating an access plan for the database query using the determined retrieval status for the resource." As discussed above in connection with claim 3, this combination of features is not disclosed or suggested by Adya. Accordingly, claims 16 and 29 are novel and non-obvious over Adya for the same reasons as presented above for claim 3. Reconsideration and allowance of independent claims 16 and 29, as well as of claims 17-28 which depend therefrom, are therefore respectfully requested.

As a final matter, Applicants traverse the Examiner's rejections of the dependent claims based upon the dependency of these claims on the aforementioned independent claims. Applicants do wish to point out, however, that a number of these claims additionally recite features that further distinguish the claims from the art of record. For example, with respect to claims 5-7 and 18-20, which recite various types of memory for which a retrieval status may be determined, the Examiner merely cites examples of those types of memory. The cited passages (e.g., in col. 5) only mention different types of memory, but disclose nothing about making determinations of whether or not certain resources are resident in those types of memories. The Examiner has therefore failed to establish that these features are disclosed by Adya.

With respect to claims 8, 13 and 21, which recite in part the determination of a percentage of a resource that has been retrieved, the cited passages are completely silent with respect to determining a percentage of a resource that has been retrieved. The Examiner has therefore failed to establish that this feature is disclosed by Adya.

With respect to claims 9-11 and 22-24, which recite in part a resource manager used to determine a retrieval status of a resource, the Examiner also relies on Levine, and in particular col. 5, line 15. However, Levine, while disclosing a "manager", discloses only a cache manager. There is no disclosure in the reference of any ability to query such a manager to determine the retrieval status of a resource, much less to do so in connection with determining costs associated with an access plan that utilizes the resource. The Examiner's rejection therefore falls far short of disclosing or suggesting each and every feature of these claims.

With respect to claims 14 and 27, which recite in part the selective generation of an access plan based upon the comparison of a current retrieval status with a stored assumption for a stored access plan, the cited passages are also completely silent with respect to this concept. The Examiner has therefore failed to establish that this feature is disclosed by Adya.

Finally, with respect to claims 15 and 28, which recite the determination of whether a beginning portion of a resource is resident in working memory, coupled with weighting a cost based upon this determination, the cited passages are completely silent

with respect to different portions of a resource, much less performing any weighting on this basis. The Examiner has therefore failed to establish that these features are disclosed by Adya.

In summary, Applicants respectfully submit that all pending claims are novel and non-obvious over the prior art of record. Reconsideration and allowance of all pending claims are therefore respectfully requested. If the Examiner has any questions regarding the foregoing, or which might otherwise further this case onto allowance, the Examiner may contact the undersigned at (513) 241-2324. Moreover, if any other charges or credits are necessary to complete this communication, please apply them to Deposit Account 23-3000.

Respectfully submitted,

June 8, 2006
Date

/Scott A. Stinebruner/
Scott A. Stinebruner
Reg. No. 38,323
WOOD, HERRON & EVANS, L.L.P.
2700 Carew Tower
441 Vine Street
Cincinnati, Ohio 45202
Telephone: (513) 241-2324
Facsimile: (513) 241-6234